Diversified Top-k Similarity Search in Large Attributed Networks

Zaiqiao Meng[a] and Hong Shen[a,b]
[a]School of Information Science and Technology, Sun Yat-Sen University, China
[b]School of Computer Science, University of Adelaide, Australia
Email: zqmeng@aliyun.com, hongsh01@gmail.com

arXiv:1608.05346v1 [cs.IR] 18 Aug 2016

***Abstract*—Given a large network and a query node, finding its top-*k* similar nodes is a primitive operation in many graph-based applications. Recently enhancing search results with diversification have received much attention. In this paper, we explore an novel problem of searching for top-*k* diversified similar nodes in attributed networks, with the motivation that modeling diversification in an attributed network should consider both the emergence of network links and the attribute features of nodes such as user profile information. We formulate this practical problem as two optimization problems: the *Attributed Coverage Diversification (ACD)* problem and the *r-Dissimilar Attributed Coverage Diversification (r-DACD)* problem. Based on the submodularity and the monotonicity of *ACD*, we propose an efficient greedy algorithm achieving a tight approximation guarantee of $1 - 1/e$. Unlike the expension based methods only considering nodes' neighborhood, *ACD* generalize the definition of diversification to nodes' own features. To capture diversification in topological structure of networks, the *r-DACD* problem introduce a dissimilarity constraint. We refer to this problem as the *Dissimilarity Constrained Non-monotone Submodular Maximization* (DCNSM) problem. We prove that there is no constant-factor approximation for DCNSM, and also present an efficient greedy algorithms achieving $1/\rho$ approximation, where $\rho \leq \Delta$, Δ is the maximum degree of its dissimilarity based graph. To the best of our knowledge, it is the first approximation algorithm for the *Submodular Maximization* problem with a distance constraint. The experimental results on real-world attributed network datasets demonstrate the effectiveness of our methods, and confirm that adding dissimilarity constraint can significantly enhance the performance of diversification.**

I. Introduction

Searching for top-k nodes similar to a given query request in a network has numerous applications including graph clustering [13], graph query [24], and object retrieval and recommendation [30]. There has been substantial research on ranking nodes and estimating the similarity (proximity) between nodes, such as the Personalized PageRank [12] and SimRank[16]. These basic methods and their variations such as P-Rank [34], TopSim [20] and Panther [33] has been successfully applied in many applications.

Nowadays with rich information available from online social networks, real social entities and their relationships can be built in a network in which nodes are associated with a set of attributes describing their properties and edges represent relationships among these nodes. In such circumstances, the problem of searching similar nodes for a given node becomes more sophisticated and challenging.

Firstly, the top-k nodes resulted from the traditional similarity searching methods are often highly related. It is hard to give a glimpse of the overall similar results with such few and highly related nodes. Search result diversification has been widely studied as a way of tackling query ambiguity and enhancing result novelty in information retrieval [5], [10]. Most of these diversified models try to trade-off the relevance for diversity in the results, thereby making the result list more diverse. In the literature there are many studies on modeling the search result diversification for network datasets, based on node's ego features diversification, such as neighborhood, exemplified by the recently proposed models of expansion ratio [21] and expanded relevance [18] applying the solution to the classic *Cardinality Constrained Monotone Submodular Maximization* problem. However, A major drawback of these models is that they have no explicit measure to eliminate redundancy resulting in a lack of novelty in their search results.

Moreover, unlike simplified structural networks only with nodes and edges, a network with attributes has more complex characters. Node attributes along with the links between them provide rich and complementary sources of information and should be used simultaneously for uncovering, understanding and exploiting the latent diversified structure in attributed network data. For example, in social networks users have profile information, and in document networks each node also contains the text of the document that it represents, hence diversification is also presence among these non-topological information. Above-mentioned diversification models can find diverse nodes which account for the topological connections between nodes, but they cannot account for the node attributes.

In this paper, to modeling the diversification search problem in attributed networks, we first formulate the problem that only considers the *attribute coverage diversification* (ACD). We prove that the optimization objective of the ACD problem is a nondecreasing submodular function, and a marginal gain based greedy algorithm can obtains an $(1-1/e)$-approximation near-optimal solution. To to improve novelty, we add an r-dissimilar constraint to ACD problem for capturing the dissimilarity between result nodes based on the topological structure of the graph. We show that the new problem become more complicated, because adding a new node to the result set following the monotonicity may break the dissimilarity constraint, the existing techniques can not be applied to our dissimilarity constrained diversification model. Our model requires to solve

a new problem of *Dissimilarity Constrained Non-monotone Submodular Maximization* (DCNSM). Based on constructing a dissimilarity-based graph, we propose a greedy algorithm achieving an approximation ratio of $1/\rho$, where ρ is bounded by the maximum degree Δ of its dissimilarity-based graph, and runs in $O((\Delta + k)(|V||A| + |E|))$ time.

The main contributions of this paper are: (1) We formalize the dissimilarity constrained diversification model for top-k diversified similarity search in attributed network that combining attribute coverage diversification with a dissimilarity constraint as the problem of maximizing a dissimilarity constrained non-monotone submodular function problem, we prove that there is no constant-factor approximation for this problem and present a linear time algorithm with approximation ratio $1/\rho$, where ρ is bounded by the maximum degree of its dissimilarity-based graph. (2) We conduct extensive experiments on real-world attributed network datasets, the results shows the effectiveness of our proposed algorithms; we also combined the baseline models with the dissimilarity constraint, and also shows that the new dissimilarity constrained methods significantly outperforms the original models.

II. Problem Formulation

A. Preliminaries

Let $G = (V, E, W)$ be an undirected weighted graph with $|V|$ nodes and $|E|$ edges. W is the *edge* weighting function such that $W = \{w : E(G) \to R+\}$. If the weighting function is not specified, the weight of each edge is 1. We denote $V(G)$ as the set of nodes in G, $G[V]$ denote the subgraph of G induced by V. Let $d_G(u)$ be the degree of node $u \in V$, Δ the maximum degree of G, $N_G(u)$ the neighborhood of v, and $N_G^+(u)$, $\{u\} \bigcup N_G(u)$.

Given an undirected weighted network G, a positive integer k, and a query node $q \in V$, the relevance metric $s(\cdot)$ is the similarity score of each nodes measuring the relevance to q. Analogously, a dissimilarity metric $diss(\cdot, \cdot)$ is the dissimilarity score measuring the dissimilarity between pair of nodes. In this paper, the similarity metric and the dissimilarity metric are defined and computed purely based on the topological structural context of nodes.

Given an element u, a ground set N, a function $f : 2^N \to R^+$ is called *submodular function* if for $\nabla S \subseteq T \subseteq N$ and $u \in N \backslash B$, $f_u(S) \geq f_u(T)$, which $f_u(S) = f(S + u) - f(S)$ is called the *marginal gain* . A submodular function f is *monotone* if for every $S \subseteq T$ we have that $f(S) \leq f(T)$ [25].

B. Problem Definition

We target the problem of diversifying top-k similarity search result for a given node in networks, assuming that the relevance score for each node the has already obtained by a given similarity search algorithms which we'll talk about more. The main challenge of this problem is how to properly measure the diversification. Some previous models try to optimize a function with a single diversification measure based on the ego features of nodes, such as neighborhood. For example, [21] formulate this problem as a bicriteria objective optimization problem that tradeoff the relevance and the *expansion ratio*, [18] optimize a single function (called *expanded relevance*) which combines both relevance and neighborhood diversity. These neighbor expansion based methods are based on the intuition that nodes with large expansion are dissimilar to each other, thus leading to diversity, and they don't have a specific metric to measure the dissimilarity ensuring the novelty. However, result nodes with large expansion set may not always indicate that they are dissimilar to each other. We illustrate this in Example 1.

Example 1. Consider a graph in Fig. 1, and suppose all the nodes has a same relevance score. The blue square nodes in Fig. 1 (a) and (b) are the possible result returned by EP1 algorithm [21] and the circle nodes denote their expansion nodes, because the expansion ratio of the selected nodes in both Fig. 1 (a) and Fig. 1 (b) is 1. But it is clearly that result nodes in Fig. 1 (b) are more novelty to each other than the result nodes in Fig. 1 (a). Because the selected nodes in Fig. 1 (b) are more dissimilar to each other whether it is measured based on shortest path or based on common neighbors.

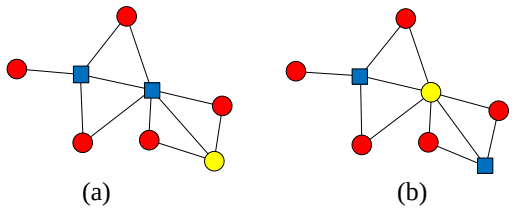

Fig. 1. (a) Diversified result from EP1 (b) A more novelty diversified result

Thus, the neighbor expansion based methods does not guarantee novelty of their result sets. Moreover, all existing algorithms do not consider the node attributes when it is involved in an attributed network.

To effectively perform diversified search, we must first determine what the diversified results is. A good diversification metric should eliminate the redundant nodes and choose a set of representative nodes which are dissimilar to each other. In real-world social networks with rich categorical node attributes, diversified search should consider both network structure and node attribute information. Our main task of this paper is to model the diversified search problem in attributed networks, considering the both network structure and node attribute information, and trading-off the diversity and the relevance.

Based on the undirected weighted graph, we give the definition of the attributed network in our problem setting.

Definition 1. Undirected Weighted Attributed Network. Let $G = (V, E, W_E, W_A)$ denotes an *undirected weighted attributed network*, where W_E is the *edge* weighting function such that $W_E = \{w : E(G) \to R+\}$, W_A

is the *node* attributed-weighting function such that $W_A = \{w : E(G) \to S_A\}$, where S_A is a family subsets of attribute set A, which representing categorical attributes of each nodes.

We assume that the attribute of nodes are binary-valued. Other types of attribute variables could be clustered into categorical variables via vector quantization, or discretized to categorical variables. For example, in Facebook social network, various universities (e.g., MIT, CMU, and Stanford) in user profile are directly treated as separate binary attributes, the posted status updates can use stemming, tokenization, and stopword removal etc. techniques to extract keywords as binary attribute. Some reduction and classification technique in attribute could be used to further improve the performance but is not the scope of this paper.

To model the diversification in attributed networks, we first introduce the *attribute coverage diversification* problem which only considering the diversification on attribute features of node, then we propose a new model of diversification combining diversification of node attributes with a dissimilarity constrain such that any two nodes in the search result must satisfy a given dissimilarity threshold to ensure the novelty, and cover attributes as much as possible.

Definition 2. ATTRIBUTE COVERAGE. Let S be a set of nodes, the *attribute coverage* of node v, $|A_v|$, is the cardinality of attribute set of v, $|A_s| = |\cup A_v|$ is the *attribute coverage* of node subset S. The *attribute coverage ratio* (ACR) of S, $ACR = \frac{|A_s|}{|A|}$, is defined as the normalization of attribute coverage, where $|A|$ is the total number of attribute set.

Problem 1. ATTRIBUTE COVERING DIVERSIFICATION (ACD). Given an undirected weighted attributed network $G = (V, E, W_E, W_A)$, a query node q, a relevant metric $s(\cdot)$, and a positive integer k, the problem is to find a subset $S \subseteq V$ that:

$$\begin{aligned} &\max_{S \subseteq V} \ f(S) = (1-\lambda) \sum_{u \in S} s(u) + \lambda \frac{|A_S|}{|A|} \\ &s.t. \ \ |S| = k. \end{aligned} \tag{1}$$

where $\lambda \in [0, 1]$ is a parameter to tradeoff relevance and diversity, $\frac{|A_s|}{|A|}$ is the *attribute coverage ratio* representing the attributed diversity. Problem 1 utilizes the linear combination of the attribute coverage and the relevant metric as the optimization objective, and aim at finding a subset S of k nodes such that: (1) the nodes in S have high relevance to the query node q; (2) the result subset S have maximum attribute coverage. The definition of attribute diversity base on attribute coverage is intuitive and reasonable, it indicate that the more attributes covered by nodes, the more diverse the result set is. However, the ACD problem only considers one diversification measure that is the attribute diversification, and still can not guarantee novelty of their result sets. To consider the diversification on network topology and improve the novelty of result, the dissimilarity in topological structure between nodes is taken into account in this problem. Formally, we then model the new dissimilarity-constrained diversified search problem.

Problem 2. r-DISSIMILAR ATTRIBUTE COVERAGE DIVERSIFICATION (r-DACD). Given an undirected weighted attributed network $G = (V, E, W_E, W_A)$, a query node q, a relevant metric $s(\cdot)$, a dissimilarity metric $diss(\cdot, \cdot)$, and a positive integer k, the problem is to find a subset $S \subseteq V$ that:

$$\begin{aligned} &\max_{S \subseteq V} \ f(S) = (1-\lambda) \sum_{u \in S} s(u) + \lambda \frac{|A_S|}{|A|} \\ &s.t. \ \ |S| = k, \\ &\qquad \forall v_1, v_2 \in S, \ \ diss(v_1, v_2) \leq r. \end{aligned} \tag{2}$$

In this problem, the r-dissimilar constraint scatters the result to a wider range of topological structure space, describes the topological structural diversification of the result set. The relevant measurement and dissimilar metrics in Problem 1 and Problem 2 will be discussed in the next section.

C. Connection to Neighbor Expansion based Methods

There exists many other previous work studies the bicriteria optimization measures in diversified search on graphs problems. For example, [21] tries to optimize the following diversified bicriteria optimization objective:

$$\max_{S \subseteq V} f(S) = (1-\lambda) \sum_{u \in S} s(u) + \lambda \frac{|N_S|}{|N|}. \tag{3}$$

The first term is the sum of the Personalized PageRank scores over the results, which reflects the relevance. The second term is the expansion ratio of the results, which reflects the diversity.

We notice that our ACD problem is the ego feature generalization of this neighbor expansion diversification problem. If we view the neighbor sets of nodes as their attributed sets, the neighbor expansion based diversified problem will become our ACD problem. In other words, the neighbors is the attributes of nodes, to some degree.

Another neighbor expansion based measure called *expanded relevance* which combines both relevance and diversity into a single function in order to measure the coverage of the relevant part of the graph [18]. They argue that bicriteria optimization is inappropriate, because the diversification methods that seem to optimize both criteria are highly correlated among each other. It is worth mentioning that both ACD and r-DACD also optimize a bicriteria objective, the structural similarity and the attributed diversity, but they obviously are two independent metrics to measure relevance and diversity. To capture the novelty among result nodes, we also add the same dissimilarity constraint to these neighbor expansion based methods. We will illustrate the comparison of these dissimilarity constrained models in detail in the follow-up experiments.

III. ALGORITHMS

A. Relevance Metric and Dissimilarity Metric

Given a large graph, finding top-k most similarity (proximity) nodes to a given query node is a fundamental problem. Most of previous work about diversified search on graphs utilize Personalized PageRank (PPR) [12] as their relevance metric. In this paper we use Panther [33] to measure the

relevance to query node and the dissimilarity between pair of nodes in result set, the major reasons are as follows. Firstly, compared with PPR, Panther has a lower time complexity which is unrelated to the node size, so that it can deal with large-scale networks. Secondly, we need an unified measure to accommodate both the relevant and the dissimilarity of nodes, and using Panther can easily meet this needs.

The basic idea of Panther is that two nodes are similar if they frequently appear on the same paths. The algorithm randomly select a vertex in G as the starting point, and then conduct random walks of T steps from v using the weight proportion as the transition probability to corresponding nodes. The relevance score is defined as: $s(u) = \frac{|p_{u,q}|}{R}$, where $|p_{u,q}|$ is the number of paths that contains node u and the query node q, and R is the total number of random paths. The authors show that Panther can provably and accurately estimates the similarity between any pair of nodes [33].

In this paper, we also evaluate the dissimilarity of nodes by the idea that two nodes are dissimilar if they don't occurs frequently together on the same paths. The normalized dissimilar score is defined as: $diss(u, v) = 1 - \frac{|p_{u,v}|}{|p_{max}|-|p_{min}|}$, where $|p_{max}|$, $|p_{min}|$ represent the maximum and minimum number of random paths between two nodes, and $|p_{u,v}|$ is the number of random paths between node u and node v.

B. The Greedy ACD Algorithm

It is easy to see that the ACD problem is NP-hard, because if we let $\lambda = 1$, the problem is equivalent to the *max k-cover* problem which is known to be NP-hard [8]. Thus we resort to develop approximate algorithms for this problem. Below, we prove that Problem 1 is a nondecreasing submodular function with a cardinality constraint [3].

Theorem 1. *The optimization problem of $f(S)$ defined in Problem 1 is a Cardinality Constrained Monotone Submodular Maximization problem.*

Theorem 1 can be proved by the same way as the proof of Theorem 3.2 in [21], so we omit the proof here.

Although maximization of submodular functions is NP-hard, there exists efficient algorithms with theoretical approximation guarantee for this problem. We propose to construct a greedy solution that repeatedly chooses a node with the *maximum marginal gain* until the size of result set equal to k. Algorithm 1 outline the *greedy ACD* algorithm (GACD).

In this algorithm, the most of time cost is computing the marginal gain $f_u(S)$ for every nodes (step 3). We can use an indicator array to represent the attribute coverage, then the time complexity from step 3 to step 4 is $O(|V||A|)$. Thus, the total time complexity is $O(k|V||A|)$.

Theorem 2. *Algorithm 1 is an $(1 - 1/e)$-approximation algorithm for Problem 1.*

According to the results in [25], for a monotone submodular function, greedily constructing by selecting an element with the maximum marginal gain gives an $(1-1/e)$-approximation to the optimal. Since from Theorem 1, we can easily prove that Algorithm 1 is an $(1-1/e)$-approximation algorithm (We omit the proof here).

Algorithm 1 Greedy ACD

Input: An undirected weighted attributed network G, k, λ, query node q, relevance metric $s(\cdot)$ to q
Output: A set S with k nodes
1: $i \leftarrow 0$, $S \leftarrow \emptyset$, $G_i = G$;
2: **while** $|S| < k$ **do**
3: update $f_u(S)$ for nodes in G_i;
4: $v \leftarrow \max_{u \in V(G_i)} f_u(S)$;
5: $S = S \bigcup \{v\}$;
6: $G_{i+1} = G[V(G_i) - v]$;
7: $i = i + 1$;
8: **return** S;

C. The Greedy r-DACD Algorithm

The main difficulty in designing algorithm for r-DACD problem is that considering the dissimilarity constraint makes the optimization objective become non-monotone, because adding a new node to the result set may break the dissimilarity constraint. The greedy algorithm used to solve the *monotone submodular maximization* problem can not directly be deployed for r-DACD problem. As far as we know, there were no reports and literature of the relevant studies mention this dissimilarity constrained problem. We denote this problem as *Dissimilarity (Distance) Constrained Non-monotone Submodular Maximization* (DCNSM) problem. In fact, the follow theorem shows that the DCNSM problem is NP-hard even hard to approximate within any constant factor.

Theorem 3. *It is hard to approximate for the DCNSM problem within $n^{1-\epsilon}$, for any ϵ>0, unless $NP = ZPP$.*

Proof: We show that DCNSM problem is reducible to the maximum independent set (MIS) problem via a costpreserving reduction; since MIS strictly equivalent to the maximum clique problem, which is hard to approximate within $n^{1-\epsilon}$, for any $\epsilon > 0$, unless $NP = ZPP$ [11].

Given a submodular function f, a ground set N, a dissimilarity function $diss(v_1, v_2)$ for every pair of elements $v_1, v_2 \in N$. To prove the theorem, we construct a new dissimilarity-based graph $G^{'} = (V^{'}, E^{'}, f)$ defined as follows: for any element $u \in N$, u is a node of $V^{'}$ in $G^{'}$, for any pair of element that violate the dissimilarity constraint ($diss(v_1, v_2) < r$), (v_1, v_2) is an edge between v_1, v_2 in $G^{'}$. f is the node submodular function such that for any subset $S \in V^{'}$, $f(S)$ satisfy the submodularity. From this, construct a simplified instance of DCNSM problem, we set the submodular function $f(v) = 1$ for any single element, $f(S) = |S|$ for any subset $S \subseteq V^{'}$. Clearly, the construction of $G^{'}$ can be done in polynomial time. We see that any solution of MIS in $G^{'}$ is precisely the subset $I \subseteq V^{'}$ maximizing $f(I)$ and satisfying the dissimilarity constraint, and vice versa.

This completes the proof. ■

Clearly, the DCNSM problem with cardinality constraint also satisfy Theorem 3. To solve the r-DACD problem is to solve cardinality constrained DCNSM problem, we put forward the *Greedy r-DACD* algorithm (GrDACD). Details of the algorithm are presented in Algorithm 2.

Algorithm 2 Greedy r-DACD

Input: An undirected weighted attributed network G, k, λ, query node q, relevance metric $s(\cdot)$ to q, a dissimilarity metric $diss(\cdot,\cdot)$
Output: A set S with k nodes
1: generate dissimilarity-based graph $G^{'}$;
2: $\rho \leftarrow 1$,$i \leftarrow 0$, $S \leftarrow \emptyset$, $G_i^{'} = G^{'}$;
3: **while** $|S| < k$ **do**
4: update $f_u(S)$ for nodes in $G_i^{'}$;
5: $v \leftarrow \max_{u \in V(G_i^{'})} f_u(S)$,
$s.t.\ f_u(S) \geq \frac{1}{\rho}\sum_{w \in N_{G_i^{'}}(u)} f_w(S)$;
6: **if** v_i is $NULL$ **then**
7: $\rho = \rho + 1$; **continue**;
8: $S = S \bigcup \{v\}$;
9: $G_{i+1}^{'} = G_i^{'}[V\left(G_i^{'}\right) - N_{G_i^{'}}^{+}(v)]$;
10: $i = i + 1$;
11: **return** S;

In Algorithm 2, the algorithm first constructing a dissimilarity-based graph $G^{'}$, then greedily select a node with the maximum marginal gain that is greater than the $1/\rho$ times of the summed marginal gain of its neighbors, where ρ is denoted as *local maximal factor*. In each iteration, the algorithm compute the marginal gain $f_u(S)$ for each $u \in V\left(G_i^{'}\right)$, and chooses a node u with maximum $f_u(S)$ and satisfying the condition: $f_u(S) \geq \frac{1}{\rho}\sum_{w \in N_{G_i^{'}}(u)} f_w(S)$, what we call as the *local maximal node*. The algorithm first set $\rho = 1$, if there is no local maximal node, then let $\rho = \rho + 1$, and continue for the next iteration. We show that the algorithm can always return k local maximal nodes within $\rho \leq \Delta$, where Δ is the maximum degree of $G^{'}$.

Lemma 1. *Algorithm 2 can terminate when $\rho \leq \Delta$.*

Proof: To prove the lemma, we prove that there always exists at least one local maximal node in $G_i^{'}$ when $\rho = \Delta$. Assume that $\rho = \Delta$, and there is no local maximal node, then for each $u \in G_i^{'}$ the following inequality holds:

$f_u(S) < \frac{1}{\Delta}\sum_{w \in N_{G_i^{'}}(u)} f_w(S)$

Accumulating this inequality for each nodes in $G_i^{'}$, we can get:

$\sum_{u \in G_i^{'}} f_u(S) < \sum_{u \in G_i^{'}} \frac{1}{\Delta}\sum_{w \in N_{G_i^{'}}(u)} f_w(S)$

The number of times $f_w(S)$ appeared on the right side exactly is the degree of w, because each $f_w(S)$ on the right side represents that there exists an edge between u and w in $G_i^{'}$. Thus:

$\sum_{u \in G_i^{'}} f_u(S) < \sum_{w \in G_i^{'}} \frac{d(w)}{\Delta} f_w(S)$

Rearranging yields:

$\sum_{u \in G_i^{'}} \frac{\Delta}{\Delta} f_u(S) < \sum_{u \in G_i^{'}} \frac{d(u)}{\Delta} f_u(S)$

From $d(u) \leq \Delta$, we reach a contradiction, hence $\rho \leq \Delta$. ■

Theorem 4. *Algorithm 2 gives a $\frac{1}{\rho}$-approximation for Problem 2, where $\rho \leq \Delta$, Δ is the maximum degree of its dissimilarity-based graph .*

Proof: We just need to prove it based on the approximate solution of the generated dissimilarity-based graph $G^{'}$. Let $OPT(G^{'})$ be optimal value for $G^{'}$, $v_i \in S$ be the selected node with the order i by Algorithm 2. The algorithm starts with $S = \{\emptyset\}$, and ends when $|S| = k$ with an optimal value of $\sum_{i=1}^{k} f_{v_i}(S)$.

Now we consider the induced subgraph $G^{'}\left(N_{G_i^{'}}^{+}(v_i)\right)$. The optimal solution of $G^{'}\left(N_{G_i^{'}}^{+}(v_i)\right)$ can be v_i or a subset of $N_{G_i^{'}}(v_i)$, and the optimal value satisfies:

$OPT\left(G^{'}\left(N_{G_i^{'}}^{+}(v_i)\right)\right) \leq \max\left(f_{v_i}(S), f_{N_{G_i^{'}}(v_i)}(S)\right)$

From the local maximal condition, we have:

$f_{v_i}(S) \geq \frac{1}{\rho}\sum_{w \in N_{G_i^{'}}(v_i)} f_w(S)$

Combining the two inequation we have:

$f_{v_i}(S) \geq \frac{1}{\rho} OPT\left(G^{'}\left(N_{G_i^{'}}^{+}(v_i)\right)\right)$

For each $v_i \in S$, their the induced subgraph $N_{G_i^{'}}^{+}(v_i)$ is independent, according to the submodularity of f, the following inequality holds:

$\sum_{i=1}^{k} OPT\left(G\left(N_{G_i^{'}}^{+}(v_i)\right)\right) \geq OPT\left(G^{'}\right)$

Then we have:

$\sum_{i=1}^{k} f_{v_i}(S) \geq \sum_{i=1}^{k} \frac{1}{\rho} OPT\left(G^{'}\left(N_{G_i^{'}}^{+}(v_i)\right)\right) \geq \frac{OPT\left(G^{'}\right)}{\rho}$

From Lemma 1, the upper bound of ρ is Δ, then we complete the proof. ■

Now we analysis the time complexity of Algorithm 2. For each loop, computing the marginal gain $f_u(S)$ for every nodes in G_i (step 4) needs $O(|V||A|)$ time. Finding the *local maximal node* in step 6 will takes $O(|E|)$ time. The loop will be terminated in at most $\Delta + k$ iterations, hence the complexity of GrACD is $O((\Delta + k)(|V||A| + |E|))$, which is linear w.r.t. the node size, the edge size and the attribute size.

Although GrDACD does not achieve any constant-factor approximation, its implements on the real-world data to conduct the experiment can achieve good results.

IV. Experimental Setup

A. Datasets

For our experiments, we consider three real-world datasets where we have network topological information as well as node attributes. The brief statistical information of our datasets are presented in Table I.

The Facebook dataset is downloaded from SNAP[1]. This dataset is built from profile and relation data from 10 users'

[1]http://snap.stanford.edu/data/

ego-networks in Facebook, and the attributes are constructed by their user profiles. The DBLP dataset is from the DBLP public bibliography data[2]. We build a coauthor network extracting from the papers in top 172 conferences (rank A and B) from 10 research areas ranked and classified by CCF[3]. We treat each author as a vertex, each collaboration of paper as an edge. The 172 conferences are viewed as the attributes of each node. The AMiner coauthor dataset is downloaded from AMiner.org[4]. This network dataset is also built by collaboration relationships among the authors, but unlike the DBLP dataset, the attributes in AMiner dataset are constructed by the extracted keyterms of their papers.

TABLE I
SUMMARY OF THE REAL-WORLD DATASETS

datasets	nodes	edges	attributes
Facebook	4,039	88,234	1,406
DBLP	73,242	373,797	172
AMiner	1,560,640	4,258,946	2,868,034

B. Evaluation Metrics

In the literature, there are no well accepted measures for diversified search in graph, since it is different to the "diversification" definition. In our experiments, we first employ two common metrics appeared in references to measure the relevance and the diversity in topological structure. The first one is the *normalized relevance* (Rel) which is given in [28], defined as $Rel = \frac{\sum_{u \in S} s(u)}{\sum_{u \in \tilde{S}} s(u)}$, where S denotes the top-k diversified result list by the diversified algorithms, $\tilde{S}$ denotes the top-k similarity nodes list by relevance metric. By definition, the normalized relevance measures the similarity of result to the query request, and the higher Rel implies the better relevance. The second metric is the *density* of the induced subgraph of the result set. The density of a graph is the number of edges existing in the graph divided by the maximum possible number of edges in the graph. In the topological structure perspective, the less density implies the more diverse of the result.

We also introduce two new metrics for our new problems: the *attributed coverage ratio* (ACR) and the *minimum dissimilarity* ($MinDiss$). The attributed coverage ratio is defined in Definition 2, which measure the attributed diversity. The average dissimilarity is defined as $MinDiss = Min_{v_1,v_2 \in S} diss(v_1, v_2)$, which measure the structural diversity.

C. Baselines

We compare our proposed methods with several state-of-the-art baselines: the bicriteria expansion optimization [21] (denoted by EP) and the expanded relevance [18] (denoted by BC). For our experiments, we mainly focus on l-step, for $l = 1$ and $l = 2$, denoted by EP1, EP2 and BC1, BC2 respectively. As EP1, EP2, BC1 and BC2 also can be build on the dissimilarity constraint version, we also consider four r-Dissimilar constraint variants of these methods, denoted by r-DEP1, r-DEP2 and r-DBC1, r-DBC2 respectively.

D. Parameter Settings and Experimental Environment

For the experiments based on diversified search problems, we fist run the Panther algorithm to obtain the relevance score for all nodes, extract the candidate nodes that have some relevance to the query node and compute the dissimilarity between all pairs of candidate nodes, then run the diversified algorithms to obtain their diversified search results. There are two parameters in Panther: path length T and error-bound ϵ. We empirically set $T = 5$ and $\epsilon \approx \sqrt{1/|E|}$ which can obtain a good performance on all the datasets [33]. In GACD, GrDACD, EP1, EP2, r-DEP1 and r-DEP2 algorithms, the candidate set size is 2000, because the maximum result set size k of our experiments is 100, and nodes with relevance rank more than 2000 are often irrelevant to the query node. In BC1, BC2, r-DBC1 and r-DBC2, the candidate nodes are those nodes with a relevance score of more than 0.0001. In diversified search algorithms, there are two parameters: λ used to tradeoff relevance and diversity, and r used to restrict the dissimilarity of results. We use the different parameter λ value in different algorithms, because the scale coefficient of their diversity function is different.

All experiments were conducted on an Ubuntu 14.04 server with two Intel Xeon E5-2683 v3 (2.0GHz) CPU and 128G RAM. All algorithms are implemented by C++.

V. RESULTS AND ANALYSIS

A. Comparison of Metrics

In this subsection, we evaluate the effectiveness of GACD and GrDACD under diversity and relevance metrics defined above. We run all the algorithms on the three real-world network datasets with varying k values ($k \in [5, 100]$). Normalized relevance (Rel) plots in the first column of Fig. 2 shows that GACD achieve comparable score in Facebook and AMiner datasets, but are worse than EP1 and EP2 in DBLP dataset. We can also observe the fact in the first column of Fig. 2 that, when we combine the dissimilarity constraint, Rel decreased significantly in all the datasets, and our GrDACD can keep a relatively high Rel score. Although a low Rel score is not an indication of being dissimilar to the query, but excessive high Rel score usually implies that the algorithm ignored the diversity of the results.

We also plot the density results in the second column of Fig. 2 and observe a similar result. When we add a dissimilarity constraint to these methods, the density decreased significantly compared with GACD and original baselines, and r-DBC gets the lowest density in all the datasets. Recall that the lower density indicates the less similar to each other, which represents the more diverse in structural topology to some extent. However, the normalized relevance and the density can not measure the diversity in attributed coverage.

[2]http://dblp.uni-trier.de/xml/
[3]http://www.ccf.org.cn/sites/paiming/2015ccfmulu.pdf
[4]https://aminer.org/AMinerNetwork

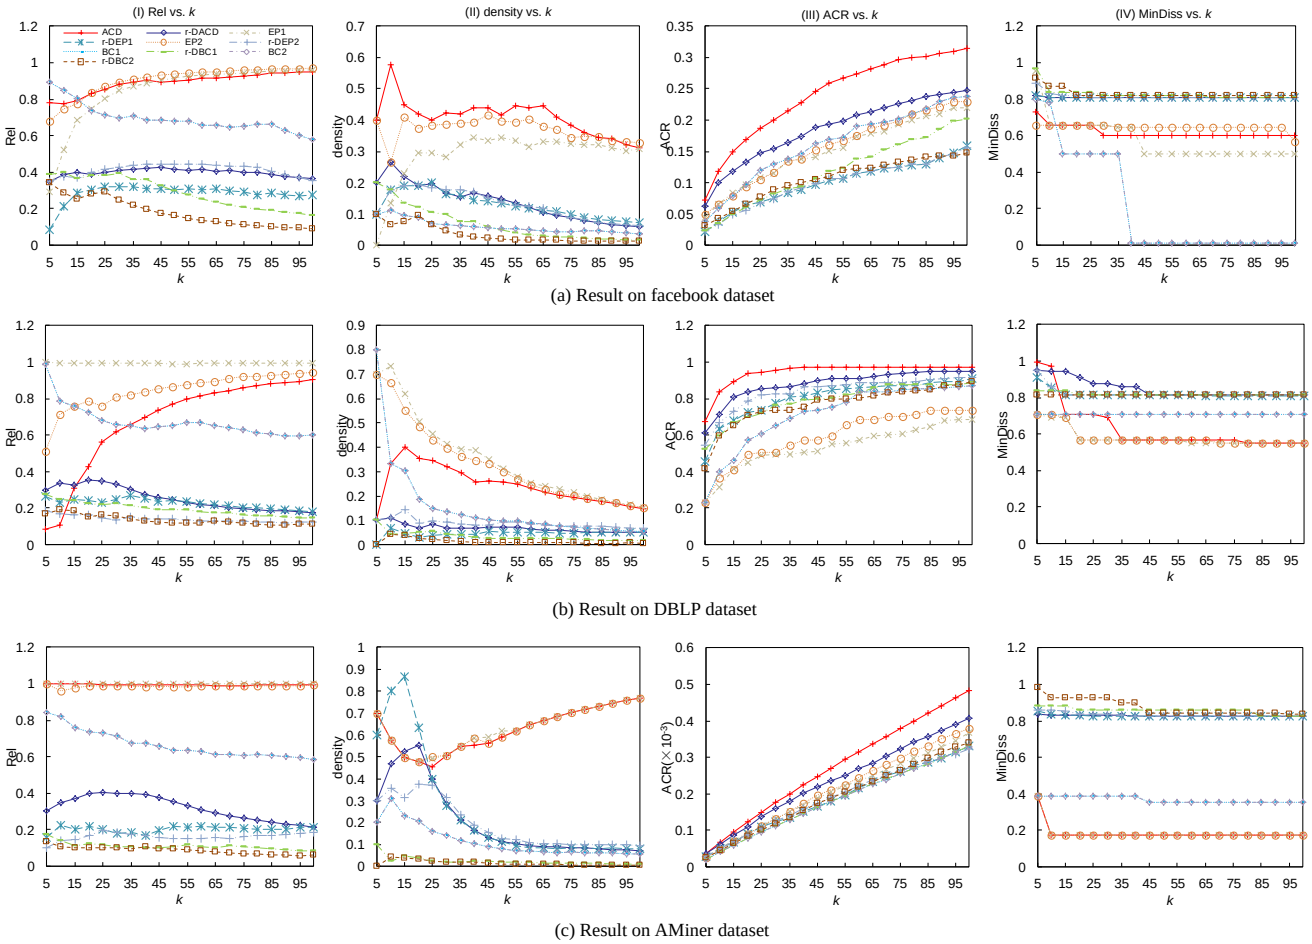

(a) Result on facebook dataset

(b) Result on DBLP dataset

(c) Result on AMiner dataset

Fig. 2. Comparison under various diversity and relevance metrics

Hence, we measure the ACR score of these returned results. Experimental comparison in the third column of Fig. 2 shows that our proposed GACD outperforms the competitors (recall that the higher value of ACR, the more diverse attributes are), and GrDACD also achieve a good score of ACR only behind GACD in all datasets. Certainly, the reason behind this result is that our GACD and GrDACD assign the attribute covering ratio as the optimizing objective, while others do not consider attributes which leading to relatively poor results. Based on the results on real-world network datasets, we conclude that our GACD and GrDACD can effectively obtain the attributed diversified search result, and the GrDACD exhibit both better attributed diversity and better structural topology diversity.

To further validate the effectiveness of the dissimilarity constraint, we evaluate the $MinDiss$ performence for all the algorithms. The results are depicted in the fourth column of Fig. 2. In this set of figures, we can clearly see that results from the dissimilarity constrained algorithms all have a higher $MinDiss$ score than GACD and the original of neighbor expansion based algorithms. And the score of BC1 and BC2 are even near zero in Facebook network, which means that there exist at least two nodes that are extremely similar to each other in the result set. We can say that those result sets are lack of novelty, and adding the dissimilarity constraint can significantly enhance the novelty.

B. Scalability

To evaluate the scalability performance of our proposed algorithms, we execute the two algorithms on a series of synthetic datasets generated by Erdös-Rényi (ER) random network model [7] over different k, with node size ranging from 50,000 to 300,000, total set-weight length ranging from 2,0000 to 120,000. With the runtime experiments shown in Fig. 3, we can clearly see that both GACD and GrACD scale linearly w.r.t the result size, the number of nodes and the number of attributes, and the GACD algorithm outperform slightly. This confirms our time complexity analysis in the previous sections, thus they all can be scalable to large networks.

C. Case Study

Now we present a case study to demonstrate the effectiveness of the proposed methods. Table II shows an example of top-5 diversified search results for Jiawei Han in DBLP network. We also count the number of conferences (attributes)

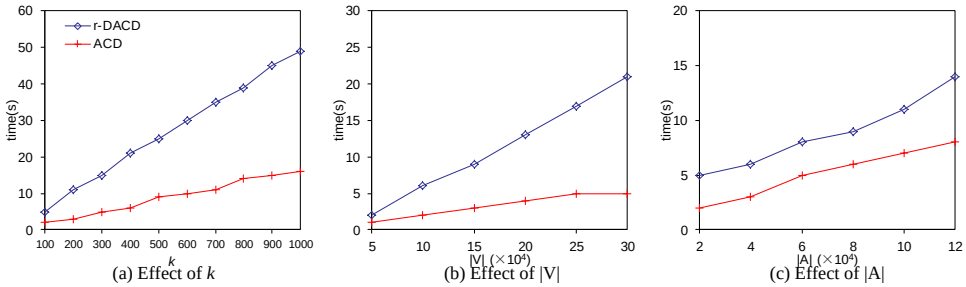

(a) Effect of k (b) Effect of |V| (c) Effect of |A|

Fig. 3. Scalability of our algorithms

TABLE II
CASE STUDY

method	Panther	ACD	r-DACD	EP1	r-DEP1	BC1	r-DBC1
TOP 5	Chi Wang 0001 Yizhou Sun Tarek F. Abdelzaher Xifeng Yan Heng Ji	Li Zhang Yizhou Sun Chi Wang 0001 Xiangyu Zhang Tarek F. Abdelzaher	Tarek F. Abdelzaher Yizhou Sun Chi Wang 0001 Li Zhang Chao Zhang	Chi Wang 0001 Tarek F. Abdelzaher Xifeng Yan Heng Ji Philip S. Yu	Xin Jin Hong Cheng Dong Wang Chao Zhang Bolin Ding	Xifeng Yan Tarek F. Abdelzaher Yizhou Sun Jing Gao Jian Pei	Philip S. Yu Thomas S. Huang Xifeng Yan Wei Liu Dong Wang
conferences	40	79	73	48	65	44	68
areas	8	10	10	8	9	8	9

that the top-5 authors covered and the number of research areas that conferences belong to. As shows in Table II, the 7 methods present very different results. Those authors found by Panther have closest relevance with the query author, but have the worst coverage in conference and research area. On the contrary, our GACD and GrDACD achieves the best and the next best performance in both conference and research area coverage. And both r-DEP1 and r-DBC1 also are significantly better performance than EP1 and BC1. These results not only confirm the effectiveness of the proposed our GACD and GrDACD methods, and also confirm that combining dissimilarity constraint can enhance the performance of diversification.

VI. Related Work

There has been various measures to estimate similarity between nodes on networks. Personalize PageRank (PPR) [12] is a random walk based measures evolved from the classic PageRank algorithm [26]. Similar to PPR, SimRank is defined recursively with respect to the "random surfer-pairs model", it evaluates the similarity between two nodes as the first-meeting probability of two random surfers. Existing algorithms like P-Rank [34], TopSim [20] are the extension of SimRank. Some other examples include discounted/truncated hitting time [27], penalized hitting probability [32], and nearest neighbor [1], [31] are also referred the random walk method. Recently, a random path sampling based method—Panther [33] was proposed that can provably, fast and accurately estimates the similarity between nodes.

Submodularity is a property of set functions with deep theoretical consequences and far-reaching applications. Submodular set functions has been widely applied to many fields, including document summarization [22], image segmentation [15], sensor placement [17], diversifying search [21], [2], and algorithmic game theory [6]. Submodular function maximization captures classic NP-hard problems in the combinational optimization such as *max cut* problems, *maximum facility location* problems and *max k-cover* problems [3], [4], [9], [29]. Some research deals with maximizing submodular functions subject to various combinatorial constraints, such as the bases of a matroid [29], multiple knapsack constraints [19] and submodular knapsack [14]. To the best of our knowledge, there is no published work providing studies in the problem of maximizing submodular functions subject to a dissimilarity (distance) constraint.

There are several studies on search results diversification in network data. DivRank[23] employs a time-variant random walk process to facilitates the rich-gets-richer mechanism in node ranking. Tong, et al. [28] propose a scalable diversified ranking algorithm by optimizing a predefined diversified goodness measure. Recently a neighbor expansion based diversified ranking method was proposed, with the assumption that nodes with large expansion would be dissimilar to each other [21] . Küçüktunç [18] propose a measure called expanded relevance which combines both relevance and diversity into a single function in order to measure the coverage of the relevant part of the graph. These methods are designed to work with the simplified structural network without considering node attributes. Our diversification models focus on the attributed networks, and combines with a dissimilarity constraint as an explicit measure to eliminate redundancy .

VII. Conclusion

In this paper, we explore a practical problem of diversifying search results in attributed networks. Based on modeling attributed diversification problem (ACD), we formulate this

problem as the r-DACD problem that combining attributed diversification with dissimilarity-constrained diversification to improve novelty of search results. We show that the r-DACD problem is hard to approximate within any constant factor. Two approximation algorithms was proposed to solve these two problems with bounded approximation ratio. We empirically compare our algorithms with two state-of-the-art diversified search methods, as well as their improved algorithms combining a dissimilarity constraint, in terms of both the structural diversified metrics and the attributed diversified metrics in real-world attributed networks. The results shows the effectiveness of our proposed algorithms, and also confirms that combining dissimilarity constraint in diversification can significantly improve the query result on novelty.

REFERENCES

[1] P. Bogdanov and A. Singh, "Accurate and scalable nearest neighbors in large networks based on effective importance," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. ACM, 2013, pp. 1009–1018.

[2] A. Borodin, H. C. Lee, and Y. Ye, "Max-sum diversification, monotone submodular functions and dynamic updates," in *Proceedings of the 31st symposium on Principles of Database Systems*. ACM, 2012, pp. 155–166.

[3] N. Buchbinder, M. Feldman, J. S. Naor, and R. Schwartz, "Submodular maximization with cardinality constraints," in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2014, pp. 1433–1452.

[4] N. Buchbinder, M. Feldman, J. Seffi, and R. Schwartz, "A tight linear time (1/2)-approximation for unconstrained submodular maximization," *SIAM Journal on Computing*, vol. 44, no. 5, pp. 1384–1402, 2015.

[5] J. Carbonell and J. Goldstein, "The use of mmr, diversity-based reranking for reordering documents and producing summaries," in *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 1998, pp. 335–336.

[6] S. Dughmi, T. Roughgarden, and M. Sundararajan, "Revenue submodularity," in *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 2009, pp. 243–252.

[7] P. ERDdS and A. R&WI, "On random graphs i," *Publ. Math. Debrecen*, vol. 6, pp. 290–297, 1959.

[8] U. Feige, "A threshold of ln n for approximating set cover," *Journal of the ACM (JACM)*, vol. 45, no. 4, pp. 634–652, 1998.

[9] U. Feige, V. S. Mirrokni, and J. Vondrak, "Maximizing non-monotone submodular functions," *SIAM Journal on Computing*, vol. 40, no. 4, pp. 1133–1153, 2011.

[10] S. Gollapudi and A. Sharma, "An axiomatic approach for result diversification," in *Proceedings of the 18th international conference on World wide web*. ACM, 2009, pp. 381–390.

[11] J. Håstad, "Clique is hard to approximate within $n^{1-\epsilon}$," in *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on*. IEEE, 1996, pp. 627–636.

[12] T. H. Haveliwala, "Topic-sensitive pagerank," in *Proceedings of the 11th international conference on World Wide Web*. ACM, 2002, pp. 517–526.

[13] X. Huang, H. Cheng, and J. X. Yu, "Dense community detection in multi-valued attributed networks," *Information Sciences*, vol. 314, pp. 77–99, 2015.

[14] R. K. Iyer and J. A. Bilmes, "Submodular optimization with submodular cover and submodular knapsack constraints," in *Advances in Neural Information Processing Systems*, 2013, pp. 2436–2444.

[15] S. Jegelka and J. Bilmes, "Submodularity beyond submodular energies: coupling edges in graph cuts," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 2011, pp. 1897–1904.

[16] G. Jeh and J. Widom, "Simrank: a measure of structural-context similarity," in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2002, pp. 538–543.

[17] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *The Journal of Machine Learning Research*, vol. 9, pp. 235–284, 2008.

[18] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Diversified recommendation on graphs: pitfalls, measures, and algorithms," in *Proceedings of the 22nd international conference on World Wide Web*. ACM, 2013, pp. 715–726.

[19] A. Kulik, H. Shachnai, and T. Tamir, "Maximizing submodular set functions subject to multiple linear constraints," in *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2009, pp. 545–554.

[20] P. Lee, L. V. Lakshmanan, and J. X. Yu, "On top-k structural similarity search," in *Data Engineering (ICDE), 2012 IEEE 28th International Conference on*. IEEE, 2012, pp. 774–785.

[21] R.-H. Li and J. X. Yu, "Scalable diversified ranking on large graphs," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 25, no. 9, pp. 2133–2146, 2013.

[22] H. Lin and J. Bilmes, "Multi-document summarization via budgeted maximization of submodular functions," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2010, pp. 912–920.

[23] Q. Mei, J. Guo, and D. Radev, "Divrank: the interplay of prestige and diversity in information networks," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. Acm, 2010, pp. 1009–1018.

[24] D. Mottin, F. Bonchi, and F. Gullo, "Graph query reformulation with diversity," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 825–834.

[25] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functionsâĂŤi," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.

[26] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking: bringing order to the web." 1999.

[27] P. Sarkar and A. W. Moore, "Fast nearest-neighbor search in disk-resident graphs," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2010, pp. 513–522.

[28] H. Tong, J. He, Z. Wen, R. Konuru, and C.-Y. Lin, "Diversified ranking on large graphs: an optimization viewpoint," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2011, pp. 1028–1036.

[29] J. Vondrák, "Symmetry and approximability of submodular maximization problems," *SIAM Journal on Computing*, vol. 42, no. 1, pp. 265–304, 2013.

[30] Z. Wang, J. Liao, Q. Cao, H. Qi, and Z. Wang, "Friendbook: a semantic-based friend recommendation system for social networks," *Mobile Computing, IEEE Transactions on*, vol. 14, no. 3, pp. 538–551, 2015.

[31] Y. Wu, R. Jin, and X. Zhang, "Fast and unified local search for random walk based k-nearest-neighbor query in large graphs," in *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 2014, pp. 1139–1150.

[32] C. Zhang, L. Shou, K. Chen, G. Chen, and Y. Bei, "Evaluating geo-social influence in location-based social networks," in *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM, 2012, pp. 1442–1451.

[33] J. Zhang, J. Tang, C. Ma, H. Tong, Y. Jing, and J. Li, "Panther: Fast top-k similarity search on large networks," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 1445–1454.

[34] P. Zhao, J. Han, and Y. Sun, "P-rank: a comprehensive structural similarity measure over information networks," in *Proceedings of the 18th ACM conference on Information and knowledge management*. ACM, 2009, pp. 553–562.